

UNIVERSITY OF
ILLINOIS LIBRARY
AT URBANA-CHAMPAIGN
BOOKSTACKS

CENTRAL CIRCULATION BOOKSTACKS

The person charging this material is responsible for its renewal or its return to the library from which it was borrowed on or before the **Latest Date** stamped below. **You may be charged a minimum fee of \$75.00 for each lost book.**

Theft, mutilation, and underlining of books are reasons for disciplinary action and may result in dismissal from the University.

TO RENEW CALL TELEPHONE CENTER, 333-8400

UNIVERSITY OF ILLINOIS LIBRARY AT URBANA-CHAMPAIGN

SEP 04 1997

DEC 10 1997

When renewing by phone, write new due date below
previous due date.

L162



Digitized by the Internet Archive
in 2011 with funding from
University of Illinois Urbana-Champaign

<http://www.archive.org/details/examinationofgen758cave>

330
B385
no. 758
Cop. 2



EBR

**FACULTY WORKING
PAPER NO. 758**

An Examination of the General Concept of Perfection

Jonathan A. K. Cave

UNIVERSITY OF ILLINOIS, URBANA-CHAMPAIGN

College of Commerce and Business Administration
Bureau of Economic and Business Research
University of Illinois, Urbana-Champaign

C 315e

BEBR

FACULTY WORKING PAPER NO. 758

College of Commerce and Business Administration

University of Illinois at Urbana-Champaign

March 1981

An Examination of the General Concept of Perfection

Jonathan A. K. Cave, Assistant Professor
Department of Economics

LIBRARY U. OF I. URBANA-CHAMPAIGN

Abstract

Among the concepts of perfect equilibrium used in the literature are: i) subgame perfection; ii) trembling-hand perfection; iii) properness; and iv) perfection. In this note, they are interpreted, compared, and extended to: games with chance players; arbitrary notions of "large mistakes"; and in the case of perfection, games without perfect recall.

An Examination of the General Concept of Perfection

For several years, game theorists have been concerned that the Nash Equilibrium solution concept may be at once too strong and too weak to provide a satisfactory model of behavior in noncooperative games. It has been accused of being too strong on the grounds that it relies on an implicit model of how a particular game is played that has very little to do with the game. In particular, the rationality assumptions underlying the concept are quite strong, and provide no room for "insurance" that players may wish to make against each other. On the other hand, in many games the set of Nash Equilibria is too large to be of much use, as in the supergame without discounting, where any feasible payoff that gives each player his minmax level (or more) is an equilibrium payoff. These are not really contradictory criticisms, since the strength of the assumptions required to support equilibrium behavior does not seem quite so extreme when there is an unique equilibrium.

More focused versions of these general criticisms have concentrated on the structure of equilibrium strategy combinations. For example, in the context of an extensive-form game, it is possible for equilibrium behavior to be secured by wholly incredible threats. Also, equilibrium behavior has only the weakest stability properties: while each person is content to follow the equilibrium prescription as long as the others do so, there is no guarantee that a small deviation on the part of one of the players, whether by accident or design, will not be followed by large deviations on the part of all the others.

To address these and similar criticisms, a number of refinements of the equilibrium concept have been developed. The origins of these

refined versions, and even their recent history, are quite diffuse, and they have gone by different names. However, perhaps the best-known line of development is that initiated by Selten in a series of papers which provide the general rubric of perfection for several equilibrium-based notions with additional properties.

In this essay, we examine the substance of these modifications, and obtain extensions of two of the perfection concepts to general games. The plan is as follows: in Section I, we describe three concepts on which perfection notions are based: backwards induction; subgame perfection; and trembling-hand perfection. Section II lays out the extensive- and normal-form models we shall use, and introduces some notation. Section III discusses the perfection notions that have appeared in the literature, and Section IV introduces the extensions to general games (with chance moves and without perfect recall).

I. Intuitive Bases for Perfection

A. Backwards Induction

In one of the first major papers on the theory of games, Zermelo proved that all finite games of perfect information (with no chance moves) have optimal strategies. The method of proof was constructive, in that it provided an explicit procedure whereby such equilibria (combinations of optimal strategies) could be calculated. Starting from each endpoint, one merely worked backwards through the tree, deciding at each stage what the player whose move it was would do. Thus, if player i was the last to move along a particular play of the game, his choice can be determined by comparing the i^{th} elements of the payoff vectors associated with the endpoints that player i could have selected. For almost all games, this

will yield a unique choice, which we can use to eliminate this last move. In other words, since we know what the last player will do in each case, we can shorten the tree by one move along each branch. The penultimate player is now the last player, and we repeat the process until there are no moves left in the game. The choices we derived along the way constitute the equilibrium strategies.

This method of solution is very appealing, when it works. The reason is that it represents a calculation which each player can (and will) make in precisely the same way, so there is no need to invoke any mechanism of communication or commitment outside the game in order to implement or justify a solution. All that is required is the minimal rationality postulate that players act to maximize their own payoffs, plus the condition that this behavioral fact is common knowledge.

The problem is that it is very easy for this mechanism to fail. For example, even in a game of perfect information, it may happen that one player does not have a unique choice at a particular vertex. In other words, there may be several choices which offer the same prospective payoff, so that the player who has to choose is indifferent between the two choices. It may also be the case that the other players are not indifferent, but they cannot predict which, if any of them will be favored by the move of the indecisive player. This does not impair the existence of an equilibrium, since we can proceed with the backwards induction given any assumption about what the player will do, and we can then calculate that player's best response in terms of the strategies thus derived for the whole game. However, the player's choice at the particular vertex then takes on the nature of a threat or a promise, and is therefore somewhat less compelling.

Also, if perfect information fails, we may not be able to carry out the backwards induction for the simple reason that informational constraints on the later players may make it impossible to predict their final choices without knowledge of the earlier moves of the other players. The following two examples show these two failures of backwards induction:




Fig. 1 Failure of B.I. due to nonuniqueness of choice




Fig. 2 Failure of B.I. due to bad information

The examples are self-evident, so their explanation is left as an exercise.

B. Subgame Perfection

What we saw in the backwards induction was a situation where each player was able to predict his own and others' moves at each subsequent vertex of the tree. In order for this to happen, it was important that each player be able to act independently at each vertex. In more general

games, where perfection information does not necessarily obtain, we can nonetheless isolate certain situations, called subgames, which have (or are defined by) the property that players can select their actions in the subgame independently of those selected in the rest of the game. Under this restriction, it is sensible to ask that players employ equilibrium strategies in each subgame. This is one definition of credible behavior; a threat or promise is credible if, given what the player proposing the threat or promise knows about the reactions of the other players, this action will actually be carried out if need arises. Of course, this is not the only possible standard for credible behavior in subgames that the game will not reach, but it has the virtue of consistency; since we are assuming Nash equilibrium behavior for the game, it would seem reasonable to impose it in the subgames. The principal objection to this line of reasoning is the following: if a player reaches a subgame that should not have been reached according to the declared strategies, this must mean that one or more players failed to carry out their declared plans. Why should a player in such a situation believe that they will continue to adhere to the rest of their strategies? Equilibrium arguments tell us that it is in their best interests to do so, but if selfish interest has already proven to be a bad predictor . . .?

However, the basis of the subgame reasoning is as described: it is important that all players be able to agree on the subgames, so that each can confidently form (equilibrium-type) expectations about the behavior of the others; moreover, within each subgame, the strategies of the players must be consistent with equilibrium behavior. To have faith in such equilibria, we must again stress that all players must agree on

what the subgames are, so that they can anticipate each other's moves correctly. Finally, we note that the concept of rationality that is being invoked here is fairly strong; it has all the strength of Nash Equilibrium, where each player assumes that the other players are also going to play according to the Nash theory, plus the additional requirement that the players assume that they will all continue to adhere to these conjectures even in the face of factual evidence to the contrary. [NB: it may be felt that we have characterized the rationality requirements of Nash Equilibrium unfairly--we would beg the reader to remember that in general, players do not have the opportunity to re-evaluate their strategy choices in light of the choices of their opponents, so that Nash-rationality must be inherently prospective.]

C. Trembling Hand Perfection

Another approach to the problem of perfection is based on the idea that players may not be certain as to the other players' ability to carry out planned moves without making any mistakes. If a player believes that there is some small but positive probability that other players will make some other moves, then it behooves that player not to rely too heavily on their declared moves. This way of playing the game is therefore based on a model of small mistakes on the part of the other players. It is in some ways a more appealing intuitive concept than subgame perfection, since it involves less inference by each player about the rationality of the others: since each player uses the declared strategies of the others as the basis for his expectations, but explicitly allows for the possibility of mistakes, discovering that a mistake has in fact been made should come as less of a shock. Moreover, it is not necessary to work

out what the set of subgames are that everyone can agree upon, since each player is only interested in what the other players do, and not at all in why they do it. On the other hand, the technical aspects of the definition produce an additional restriction; while we have in mind that an equilibrium is trembling hand perfect if it is a limit of equilibria calculated for games where players are unable to expect their opponents to do exactly what they said they were going to do, the way the definition is worded suggests that this perfection can be invoked if there exists a single sequence of "models of mistakes" that has the desired property, not that convergence should hold for all such models.

One particular kind of model would admit the possibility of mistakes, but would not have them be wholly random. For example, we might require that the probability of a person making a particular kind of mistake should be proportional to (or at least an increasing function of) the profitability of the mistake. Since this requirement disallows some of the models of mistakes we were able to invoke for perfection, this leads to a strictly stronger concept, which has been called properness. Another model which seems appropriate when strategies are chosen from separable metric spaces, would restrict attention to situations where "large" mistakes are assigned lower probabilities than "small" mistakes. This is merely another way of defining the concept of properness, however, if the payoff functions are jointly continuous in all the strategies. The situations where the two concepts diverge tend to be those extensive-form games where the topology that is natural is imposed separately at each information set. In such situations, the payoff function for the game may, given the strategies of the other players, be a discontinuous

function of the choice at any particular information set, especially if mixed strategies are ruled out.

D. General Perfection

The strongest perfection concepts are an amalgam of the subgame and trembling hand ideas, although they are most easily explained in terms of subgame-type thinking. The basic idea is that each player faces several independent decision situations during a game. A decision situation is a collection of information sets which are independent of other information sets by virtue of perfect recall, and which are linked to one another by failures of perfect recall. For instance, each information set in a game of perfect recall represents an independent decision situation. The requirement is that each player, in each independent decision situation, should make a "local best reply", which is to say that each player should choose an optimum in the decision problem obtained by filling in the declared strategies of the other players, and of the particular player under consideration at other independent decision situations. There is an important difference between this notion and subgame perfection, and it comes from the fact that we no longer require all the players to share the same set of independent decision situations. This is a subtle point, and worth making carefully. When we say that a player takes the actions of all other players, including himself at other decision situations, as given, we are putting him in a conceptual framework that is closely akin to the Nash Equilibrium framework. It is not that the concerned players play each game separately, and obey simultaneity within each subgame. Rather, all subgames are considered at one time, and the selection of each part of a player's strategy is subject to the same ex post stability

requirements (i.e., best response) as the strategy as a whole. To put it slightly differently: in subgame perfection the notion of Nash Equilibrium is used repeatedly, at each subgame. For general perfection, it is used only once, but on a very disaggregated version of the game; each player being split into his independent "personalities" or agents, who are bound to each other only by the fact that they share the same payoff function.

II. Extensive- and Normal-form Models, Preliminary Definitions

We begin with the definition of a game in extensive form. We remark that these models all conform to standard usage; for further details, the reader is referred to referenced work by Kuhn and by Luce and Raiffa.

2.1 Definition: An extensive-form game is a collection $\Gamma = [K, N, \Xi, h, p]$, where:

N is the (finite) set of players;

K is the game tree, consisting of vertices and choices at each vertex;

$\Xi = [\Xi_1, \dots, \Xi_n]$ is the information structure; if we denote by K_0 those vertices of K at which the chance player moves, and by W the set of endpoints of K (vertices for which the set of choices is empty), then Ξ is a partition of $K - K_0 - W$. Each $\Xi_i = [U_i^1, \dots, U_i^t, \dots]$, where U_i^t is called an information set of player i, and consists of a set of vertices x with the following properties:

- i) the number of choices at each x is the same;
- ii) player i has to choose at each x ;

- iii) when player i chooses, he must do so in ignorance of which x he has reached; all he knows is that he is at one of the vertices in U_1^t ; and
- iv) there is no vertex x in U_1^t with the property that the (unique) path from the origin through x contains another vertex y in U_1^t .

$h:W \rightarrow R^n$ is the payoff function, and assigns an n -tuple of payoffs to each endpoint;

p is the probability structure: it describes the way chance selects moves at vertices in K_0 . For most purposes, it is sufficient to treat this as a collection of independent probability distributions on the choices at each $x \in K_0$.

2.2 Definition: Let x be a vertex of K . Since K is a tree, there is a unique path from the origin to x ; we denote this path $\bar{x}$. If $y \in \bar{x}$, we say that $y \rightarrow x$, or "y comes before x." This is a transitive relation, and partially orders the vertices of K . Now define

$$D(x) = \{y \in K: x \rightarrow y\}$$

this set is called the descendants of x . Among the vertices in $D(x)$ are vertices y with the property that there is no vertex z s.t.

$$x \rightarrow z \rightarrow y$$

these vertices are called the immediate successors of x , and denoted $I(x)$. There is a natural isomorphism between the choices at x and the members of $I(x)$; to any $y \in I(x)$ there corresponds a choice k such that if k is chosen at x , the next vertex will be y . We say that "y follows x via the k^{th} choice." We write this as $y = I(x,k)$, and define

$$D(x,k) = I(x,k) \cup D(I(x,k))$$

This is the set of vertices which follow x via the k^{th} choice. When we are dealing with sets of vertices, such as information sets, we shall write $A \rightarrow B$ if there exist $x \in A$ and $y \in B$ s.t. $x \rightarrow y$. Moreover, we define $D(A)$ to be $\bigcup_{x \in A} D(x)$, and $D(A,k) = \bigcup_{x \in A} D(x,k)$.

2.3 Definition: Let Π be a partition of Ξ_i into elements $p_1, \dots$, each of which is a collection of information sets belonging to player i . A pure p-strategy is a rule which assigns a choice to each information set in p . A pure Π -strategy is a rule which assigns a p -strategy to each $p \in \Pi$. A mixed p-strategy is a random variable whose values are pure p -strategies, while a mixed Π -strategy is a rule which assigns a mixed p-strategy to each $p \in \Pi$. Note that it follows that choices may be correlated between information sets in the same $p \in \Pi$, but must be independent between elements of Π .

2.4 Remark: This definition is not standard usage, and it may be useful to relate it to standard definitions by taking two examples:

Π is the coarse partition, it has but a single element, so a Π -strategy is a mixed strategy, in the usual terminology;

$\bar{\Pi}$ is the fine partition, each information set being a separate element, so a $\bar{\Pi}$ -strategy is a behavioral strategy.

2.5 Definition: Let Π be a partition of Ξ_i . We say that player i has perfect recall w.r.t. Π if, for any $p, p' \in \Pi$, and for any two information sets $u \in p, v \in p'$, either

- i) $v \cap D(u) = \emptyset = u \cap D(v)$, or
- ii) if (w.l.o.g.) $u \cap D(v) \neq \emptyset$, then $u \subset D(v,k)$, some k .

2.6 Remark: This also differs somewhat from standard usage; in the usual framework, an agent has perfect recall iff he has perfect recall w.r.t. $\bar{\Pi}$. It will be noted that every agent has perfect recall w.r.t. Π .

2.7 Definition: Let Γ be a game in extensive form with no chance moves [i.e., $K_0 = \emptyset$]. We say that agent i has perfect information iff every information set in Ξ_i consists of exactly one vertex. If every agent has perfect information, we say that Γ has perfect information. In a game with chance moves, we say that agent i has perfect information iff, for every information set $u_i \in \Xi_i$, and for every information set $v \in \Xi_j$ (where j may be i), we have

$$v \rightarrow u_i \Rightarrow u_i \cap D(v) \subset D(v, k)$$

for some k . An alternative, stronger requirement, is

$$v \rightarrow u_i \Rightarrow u_i \subset D(v, k)$$

for some k .

We shall now define the expected payoff function for an extensive-form game. Given any n -tuple of pure strategies for the personal [non-chance] players in Γ , which we denote by $\sigma = [\sigma_1, \dots, \sigma_n]$, and any single pattern of choices at vertices in K_0 (realization of chance moves), which we shall denote σ_0 , the game will reach a unique endpoint $w(\sigma, \sigma_0)$. Given any n -tuple of mixed strategies, we can calculate the equivalent Π -strategy, which consists of a collection of pure strategies with probabilities for each one. In exactly the same way, we can associate with any realization of the chance player's moves the probability (given by p) of its occurrence. For any such collection of choices σ_i by player i , we shall write $\text{pr}(\sigma_i)$ for the probability of its occurrence. If we wish to describe the expected payoff resulting from the use of a particular

n-tuple of mixed strategies $\mu = [\mu_1, \dots, \mu_n]$, we reflect this by writing the probabilities as $\text{pr}(\sigma_i | \mu_i)$. For the chance player, this is denoted $\text{pr}(\sigma_0 | p)$, since the chance player's strategy is fixed at p . It follows that the expected payoff is given by:

$$H(\mu) = \sum_{\sigma} h(w(\sigma_0, \sigma_1, \dots, \sigma_n)) (\text{pr}(\sigma_0 | p)) \left(\prod_{j=1}^n \text{pr}(\sigma_j | \mu_j) \right)$$

We now turn to the definition of the normal-form game.

2.8 Definition: A normal-form game is an assemblage $[N, \Sigma, h]$ where

N is the finite set of players;

$\Sigma = \Sigma_1 \times \dots \times \Sigma_n$ is the strategy space (a Cartesian product of compact sets of pure strategies); and

$h: \Sigma \rightarrow R^n$ is the payoff function.

2.9 Definition: For each player i in a normal-form game, a mixed strategy is a Σ_i -valued random variable: it induces probabilities on Σ_i , which allow us to define the expected payoff function, H , in a manner precisely analogous to the above. For each i , we denote the space of mixed strategies by M_i . As a matter of notation, whenever we refer to a Cartesian product which includes all the players except for player i , we shall use the subscript "(i)".

To every extensive form game, Γ , and partition $\Pi = \Pi_1, \dots, \Pi_n$ of the information structure Ξ , there corresponds a normal-form game $G(\Gamma, \Pi)$ defined as follows: the players in $G(\Gamma, \Pi)$ are the agents of players in Γ , where an agent is a player who controls the choice at a single element of Π_i , and has player i 's payoff. What is usually denoted "the normal form of Γ " is here denoted $G(\Gamma, \underline{\Pi})$.

2.10 Definition: A strategy μ_i of player i is said to be dominant in a set $B \subset M_i$ of player i 's strategies w.r.t. a set $A \subset M_{(i)}$ of the other players' strategies iff, for all $\mu_{(i)} \in A$ and all $\mu'_i \in M_i$,

$$H_i(\mu_i, \mu_{(i)}) \geq H_i(\mu'_i, \mu_{(i)})$$

If the above equation holds with strict inequality for all $\mu'_i \neq \mu_i$, we say that μ_i is strongly dominant in B w.r.t. A . The set of dominant strategies in B w.r.t. A is denoted $\Delta_i(B, A)$, while the set of strongly dominant strategies in B w.r.t. A is denoted $\delta_i(B, A)$ - it is either a singleton or empty.

A Nash Equilibrium is an n -tuple of strategies $\mu = [\mu_1, \dots, \mu_n]$ with the property that, for each i ,

$$\mu_i \in \Delta_i(M_i, \mu_{(i)})$$

Another way to say this condition is that μ_i is a best response to $\mu_{(i)}$. It will be noted that the set of best responses is the convex hull of the set of pure-strategy best responses.

III. Perfection Notions from the Literature

A. Pl: "Subgame Perfection" a la Selten

We begin with a general definition of a subgame defined by a collection of vertices. This definition is much more general than we need for this section, but it will be useful in Section IV, since it precisely captures the intuition behind subgame perfection.

3.1 Definition: Let Γ be an extensive-form game, and B a subset of the vertices of $K - W$. We say that B defines a subgame Γ_B iff:

- i) for each i and each u_i , $[B \cap D(u_i) \neq \emptyset] \Rightarrow [B \subset D(u_i, k)]$ some k ;
- ii) for each i and u_i , $[u_i \cap (B \cup D(B)) \neq \emptyset] \Rightarrow [u_i \subset (B \cup D(B))]$.

The first condition means that every player in or after B has perfect recall of the events leading up to B , so that no player must choose a strategy in the subgame that depends on previous choices. The second condition says that every player in the subgame is aware of being in the subgame, so that choices in the subgame can be made independently of other subgames. We shall return to this definition of a subgame later: for the moment we limit ourselves to subgames as Selten defined them, which is to say, subgames in our sense defined by single-element information sets. By the tree property, this means that condition i is trivially satisfied, and we only have to verify condition ii, which can be interpreted as the requirement that the "subtree" starting from the information set that defines the subgame does not "cut" any other information set, and if v is any other information set s.t. $u^* \rightarrow v$, we require that $v \subset D(u^*)$.

3.2 Definition: Let B define a subgame; we now define

- i) $N_B = \{i: \Xi_i \cap (B \cup D(B)) \neq \emptyset\}$; the set of players in Γ_B ;
- ii) $K_B = B \cup D(B)$; the "tree" of Γ_B (not really a tree, but a union of trees;
- iii) $\Xi_B = \{\Xi_i^B: i \in N_B\}$, where each $\Xi_i^B = K_B \cap \Xi_i$; this is well-defined by condition 3.1.ii, and is the information structure for Γ_B ;
- iv) $W_B = K_B \cap W$ - the set of endpoints of Γ_B ;
- v) h_B is the restriction of h to W_B ;
- vi) p_B is the restriction of p to $K_0 \cap K_B$;

vii) Σ_i^B is the space of player i's pure strategies in Γ_B ; the generic element σ_i^B selects a choice at each $u_i \in \Xi_i^B$. A partition of Ξ_i is said to be B-allowable if, for each $p \in \Pi$, $p \cap \Xi_i^B \neq \emptyset \Rightarrow p \subset \Xi_i^B$; of course, Π is almost never B-allowable, so we denote the coarsest B-allowable partition by Π^B --all of Ξ_i^B is a single element. Mixed Π^B -strategies have an independent component which determines the choice at this element, so we call this component a mixed strategy for Γ_B ; the space of such is denoted M_i^B .

3.3 Definition:

- i) let $\{B\}$ be a collection of subsets of K ; $\{B\}$ is proper if the origin belongs to $\{B\}$, and if each $B \in \{B\}$ defines a subgame;
- ii) since we have defined payoffs and strategies for Γ_B , it is trivial to define dominance for Γ_B , so we have the symbols: $\Delta_i^B(C,A)$ and $\delta_i^B(C,A)$ for $C \subset M_i^B$ and $A \subset M_{(i)}^B$.

We shall define subgame perfection with respect to a proper collection of subsets. Before we do so, however, we need to be certain that we are dealing with strategies which can be defined for the various subgames; since we want the definition to apply to arbitrary Π -strategies, and since Π is in general not B-allowable, we have to show that there are equivalent strategies which are Π -strategies for a B-allowable Π .

3.4 Lemma: Let Γ_B be defined, and let μ be an n-tuple of Π -strategies.

Define Π_i^B for each player i as $\Pi_i^B = [p_i^B, q_i^B]$, where

$$p_1^B = \{u \in \Xi_1 : u \in \Xi_1^B\}$$

$$q_1^B = \{\Xi_1 - \Xi_1^B\}$$

Let Π^B denote $\{\Pi_1^B, \dots, \Pi_n^B\}$. Then there exists an n -tuple of Π^B -strategies v with the property that v and μ induce the same probability distribution on W : i.e. $H(v) = H(\mu)$,

Proof: Condition 3.1.i and Kuhn's Theorem.

3.5 Definition: Let $\{B\}$ be a proper collection of subsets, and let μ be an n -tuple of Π -strategies. We say that μ is Pl-perfect w.r.t. $\{B\}$ iff, for every $B \in \{B\}$, we have

$$\mu_1^B \in \Delta_1^B(M_1^B, \mu_{(1)}^B)$$

where μ_1^B is the Γ_B part of the equivalent strategy given by the above lemma.

What this definition says is that a Pl-perfect n -tuple of strategies w.r.t. $\{B\}$ is one that induces an equilibrium in every subgame defined by a member of $\{B\}$. In case $\{B\} =$ the set of all single element information sets that define subgames, we say that the n -tuple strategies is Pl-perfect. In case $\{B\}$ is the set of all subsets that define subgames, we say the strategies are P2-perfect.

B. P3: "Trembling-hand perfection"

This perfection notion is only defined for normal-form games, so to think of it as a concept for extensive-form games requires that we pass to a normal form. We remind the reader that this definition can be applied to any normal form $G(\Gamma, \Pi)$ [Def. 2.9], but what we shall denote by P3 is the application of this definition to $G(\Gamma, \Pi)$.

3.6 Definition: Let $[N, \Sigma, h]$ be a normal form game, and represent a mixed strategy for player i by a Borel-measurable map $\mu_i: [0,1] \rightarrow \Sigma_i$. Such a map defines a measure $\lambda \circ \mu_i^{-1}$, where λ is Lebesgue measure. For convenience, given any measurable map $v: [0,1] \rightarrow \Sigma_i$, let us write $\lambda(v)$ for $\lambda \circ v^{-1}$. Such a measurable map, v , is a:

- i) mixed strategy if $\lambda(v)(\Sigma_i) = 1$; and a
- ii) minimum probability assignment if $\lambda(v)(\Sigma_i) \leq 1$.

3.7 Definition: Let v be a minimum probability assignment. It is said to be completely mixed iff there exists a Lebesgue integrable function $f: \Sigma_i \rightarrow \mathbb{R}$ s.t.

- i) $f > 0$ uniformly on Σ_i ;
- ii) for each measurable $S \subset \Sigma_i$, $\int_S f(\sigma) \lambda(v)(d\sigma) = \lambda(v)(S)$; and
- iii) $\int_{\Sigma_i} f(\sigma) \lambda(v)(d\sigma) \leq 1$.

The set of completely mixed minimum probability assignments is denoted T_i . The cartesian product $\times_{i \in N} T_i = T$.

3.8 Definition: Let $v \in T$. We define a normal-form game $[N, M(v), H]$ by $M(v) = \times_{i \in N} M_i(v_i)$ where

$$M_i(v_i) = \{\mu_i \in M_i: \text{for each measurable } S, \lambda(\mu_i)(S) \geq \lambda(v_i)(S)\}$$

If we think of this as a normal-form game, and of $M(v)$ as a space of n -tuple of pure strategies, it is fairly easy to show that the game has equilibria in pure strategies, by compactness and concavity arguments. These pure-strategy equilibria will be completely mixed strategy n -tuples for the original game, although they will not always be equilibria.

3.9 Definition: Let $[N, \Sigma, h]$ be a normal-form game, and $\mu \in M$ a Nash equilibrium. We say that μ is trembling hand perfect iff there exists a sequence v^t of members of T with the properties:

- i) $v^t \rightarrow 0$ in the weak* topology, where 0 is the trivial member of T ;
- ii) there is a sequence μ^t of members of M s.t.
 - a) $\mu^t \in M(v^t)$ for each t ;
 - b) $\mu_i^t \in \Delta_i(M_i(v_i^t), \mu_{(i)}^t)$ for each i ; and
 - c) $\mu^t \rightarrow \mu$.

We remark that, if $\text{supp}(\mu_i)$ denotes the set of pure strategies of player i (members of Σ_i) that are used with positive probability in μ_i , the following properties obtain:

- i) if μ is a Nash Equilibrium, then, for each i and each $\sigma_i \in \text{supp}(\mu_i)$, $H_i(\mu_{(i)}, \sigma_i) = H_i(\mu_{(i)}, \mu_i)$;
- ii) if μ is a Nash Equilibrium with the property that, for each i , $\text{supp}(\mu_i) = \Sigma_i$, then μ is trembling-hand perfect; and
- iii) for games with an a priori restriction to pure strategies, the concepts of equilibrium and trembling-hand perfection coincide.

One particularly simple class of games is the class of 2x2 games; those with two players, each of whom has two pure strategies. For such games, a complete classification can be obtained.

3.10 Theorem: In a 2x2 game exactly one of the following situations obtains:

- i) there is an unique equilibrium [pure or mixed] and it is trembling-hand perfect;

- ii) there are two equilibria [both pure], of which only one is trembling-hand perfect;
- iii) there are three equilibria [two pure and one mixed], all of which are trembling-hand perfect;
- iv) there are a continuum of equilibria [semi-pure: one player uses a pure strategy] of which only one, a pure one, is trembling-hand perfect.

Proof: We begin by showing that our classification of the equilibrium sets is correct. The following is a classification of all 2x2 games according to dominance:

Class I: both players have strongly dominant strategies;

Class II: one has a strongly dominant strategy; the other has a

Class III: both players have weakly dominant strategies;

Class IV: one has a strongly dominant strategy, the other has no dominant strategies;

Class V: one has a weakly dominant strategy, the other has no dominant strategies;

Class VI: neither has any dominant strategies.

The following are the equilibrium sets corresponding to each class:

Class I: unique equilibrium - each player uses his strongly dominant strategy;

Class II: there are two cases - if the two strategies of the player with the weakly dominant strategy are equivalent [give the same payoff] when the other player uses his strongly dominant strategy, then there are a continuum of equilibria, if not, there is again a unique pure-strategy equilibrium point;

Class III: there are four cases - for each player i , we denote that player's weakly dominant strategy by g_i and his other strategy by b_i .

Case III-1: $h_1(g_1, b_2) = h_1(b_1, b_2)$; $h_2(b_1, b_2) = h_2(b_1, g_2)$: here there are two isolated pure-strategy equilibrium points: (b_1, b_2) and (g_1, g_2) ;

Case III-2: $h_1(g_1, b_2) = h_1(b_1, b_2)$; $h_2(g_1, b_2) = h_2(g_1, g_2)$: here there are a continuum of equilibria of the form (g_1, q) for $q \in [0, 1]$;

Case III-3: $h_1(g_1, g_2) = h_1(b_1, g_2)$; $h_2(b_1, b_2) = h_2(b_1, g_2)$: this is symmetric with III-2, so the equilibria are points of the form (p, g_2) for $p \in [0, 1]$;

Case III-4: $h_1(b_1, g_2) = h_1(g_1, g_2)$; $h_2(g_1, b_2) = h_2(g_1, g_2)$: we get the set of points which are either of the form (p, g_2) or (g_1, q) for $p, q \in [0, 1]$.


Class IV: again, there is a unique pure equilibrium in this class; the player with the strongly dominant strategy uses it, and the other player uses his best response which is by definition of this class pure.

Class V has two cases, depending on whether the equivalence between the "b" and "g" strategies of the player with the weakly dominant ["g"] strategy occurs at a strategy of the other player that is on that player's best response function. The difference is minor, however, and in both cases there are a continuum of equilibria which will involve a mixed strategy by the player with the weakly dominant strategy. In the second case, where the equivalence of the "b" and "g" strategies occurs off the other player's best response function, there is an additional, isolated pure strategy equilibrium, where the first player plays his "g" strategy as an unique best response.


Class VI can be most easily analyzed by visualizing the best response functions of the players. By definition, if these functions are shown in the unit square, with one axis for the probability used by each player, each best-response function will connect diagonally-opposed corners [see discussion infra for pictures]. If they touch different corners, there is a unique- and completely mixed-equilibrium, while if they connect the same pair of corners there are three equilibria, one of which is mixed, while the other two are pure.

In the figures which follow, we have shown the pair of best response functions for the various cases discussed above. The segment of player i 's best response function labelled S_i is the segment where the best response is (locally - holding the other player's strategy fixed) unique. Where there are multiple best responses, they are indicated by W_i . Equilibria are given by points of intersection between the two best response functions, and we note that the following is true:


The intersection of S_1 and S_2 is always a trembling-hand perfect equilibrium point. The only other trembling-hand perfect equilibria are intersections of W_1 and W_2 that are interior to both.




I




II-1



II-2



III-1



This completes the proof, and our discussion of perfection in the P3 sense, although we shall return to the subject in the next section.

C. P4: Properness

The term "proper equilibrium" was coined by Myerson to describe a situation in which the probability of making a given mistake that each player ascribes to the others is taken to be proportional to its profitability. It is evident that, with such a restriction placed on the allowable sequences of completely mixed minimum probability assessments, the sets T_i from which they can be chosen shrink. As a consequence, the set of proper equilibria is a subset of the set of perfect equilibria. If the definition were rephrased so that perfection required convergence for every sequence of minimum probability assessments, this inclusion

would be reversed. At any rate, the intuition underlying Myerson's definition is that "small" mistakes are intuitively more likely than "large" ones. Of course, the crucial question is then; what do we mean by a small mistake?

In Myerson's work, a definition is given in terms of the payoff function: a small mistake is one that gives nearly the same payoff as what was intended.

3.11 Definition: Let $[N, \Sigma, h]$ be a normal form game, and let M_i be player i 's space of mixed strategies, with $H: M \rightarrow R^n$ being the expected payoff function. We say that $\mu \in M$ is a proper equilibrium iff there exist sequences $\varepsilon^t \subset R_+$ and $\mu^t \subset M^\circ$ (the interior of M) s.t.

- i) for each t , $\varepsilon^t > 0$, and $\lim_{t \rightarrow \infty} \varepsilon^t = 0$;
- ii) for each t , each i and each pair $\sigma_i, \sigma'_i \in \Sigma_i$, we have
 $[H_i(\sigma_i, \mu_{(-i)}^t) < H_i(\sigma'_i, \mu_{(-i)}^t)] \Rightarrow [\mu_i^t(\sigma_i) < \varepsilon^t \mu_i^t(\sigma'_i)]$; and
- iii) $\lim_{t \rightarrow \infty} \mu^t = \mu$.

In other words, the model of mistakes involved in proper equilibria is such that inferior pure strategies are given smaller and smaller weight. Since this is merely a restriction on the sequences of strategies which can be used to approach a perfect equilibrium, it is clear that this forms a subset of the set of perfect equilibria.

D. P5: Perfection

This perfection notion is a subgame perfection notion, but it strengthens the requirements of subgame perfection considerably. As originally described, it was intended to apply to games with perfect recall only. The advantage of such games is of course that mixed

strategies are equivalent to behavioral strategies, so that each information set becomes an independent decision situation. The perfection notion thus arrived at has two definitions: one is essentially that each player at each information set, taking as given the actions of all other players, and his own actions at all other information sets, is playing a "local best response". The other is that the strategies are trembling-hand perfect in the "agent-normal form" $G(\Gamma, \overline{\Pi})$.

We begin with the first definition: let $\beta \in M$ be an n -tuple of behavioral strategies for an extensive-form game Γ . To be explicit, we shall represent the strategy space of behavioral strategies as B , where

$$B = \times_{i \in N} B_i$$

$$B_i = \times_{u \in \Sigma_i} M(u)$$

and, for any information set u , $M(u)$ is the space of random decisions at u . For the generic member $\beta \in B$, we shall let $\beta(u)$ be the induced probability distribution on choices at u . Given β , we shall define a decision problem for u as follows: the alternatives are members of $M(u)$, and the payoff to each alternative is denoted $H(u|\beta)$ and calculated as follows: let $x \in u$. If the game had reached x , and the t^{th} alternative were chosen, the next vertex would be $I(x,t)$. From $I(x,t)$, the choices specified by β allow us to calculate a terminal probability distribution conditional on $I(x,t)$, which we denote $v(w,x,t)$ for an arbitrary endpoint $w \in W$. The payoff conditional on (x,t) is denoted

$\bar{h}(x, t) = \int_W h(w) v(dw, x, t)$. Of course, we do not know which element of u has been reached. However, the information in β can be used to provide a conditional distribution $v(x|\beta)$ on the members $x \in u$. This is calculated in a straightforward Bayesian manner, and this requires some comment. In general, for any vertex x and any mixed strategy $\mu \in M$, we can calculate $\text{pr}(x|\mu)$, the probability of reaching x when the players are using μ . We shall therefore write

$$v(x|\beta) = \frac{\text{pr}(x|\beta)}{\int_u \text{pr}(y|\beta) dy} = \frac{\text{pr}(x|\beta)}{\int_u \text{pr}(dy|\beta)} \quad [\text{by an alternative notation}]$$

What may be noticed is that this rule tells us nothing about what to do when $\int_u \text{pr}(y|\beta) dy = \text{pr}(u|\beta)$, the conditional probability of reaching u given β , is 0. In other words, if β has the property that some information sets are unreachable given β , the agents at those information sets do not face well defined problems.

The payoff function for the decision problem can now be written explicitly, for those agents who are "active" given β . This set of agents is denoted $A(\beta) = \{u \in \Xi: \text{pr}(u|\beta) > 0\}$.

3.12 Definition: Let $u \in \Xi$, $\beta \in B$, be given. If $u \in \Xi_1$, and if $\mu \in M(u)$ is a random choice at u , the expected payoff to μ given β is denoted $\mathcal{H}_u(\mu|\beta)$ and calculated as

$$\mathcal{H}_u(\mu|\beta) = \int_u \int_{T(u)} v(x|\beta) \mu(t) \bar{h}_1(x, t) dt dx,$$

where $T(u)$ is the set of pure choices at u .

Notice that this definition can be adapted to deal with information structures other than perfect recall; this will be done in Sec. IV.

3.13 Definition: $\beta \in B$ is perfect(P5) iff, for each $u \in \Xi$, and for each $\beta'_u \in M(u)$, where $\beta = \{\beta_u : u \in \Xi\}$, we have

$$\mathcal{H}_u(\beta_u | \beta) \geq \mathcal{H}_u(\beta'_u | \beta)$$

There is another approach to this concept of perfection. Inasmuch as the various information sets do constitute independent agents of the players, it is possible to form the agent-normal form $G(\Gamma, \overline{\Pi})$. It turns out to be the case that trembling-hand perfection in the agent-normal form is equivalent to P5 perfection [see Selten].

IV. Extensions of the Above Concepts

We have essentially three objectives in this section. First, we shall extend the notion of properness to deal with other possible topologies on the space of strategies, to see what implications can be drawn from the notion that "small" mistakes are somehow more likely than "large" mistakes. Among the points of interest here are: the different topologies that are reasonable when the normal form is derived from an extensive form; and the concept of rationality embodied in the anticipation of mistakes. The second objective is to make precise the relation between the trembling-hand and subgame perfection concepts, and to describe fully the inclusion relations between the various notions for various classes of games. Finally, we shall extend the concept of P5 perfection [optimality at each independent decision problem] by describing a canonical form of an extensive-form game that isolates the independent decision problems. There are several ways to describe this form, each of which adds something to our intuition about

the reasonableness of perfection as a refinement of the notion of Nash Equilibrium.

A. Refinements of Properness: different topologies

Since properness derives from trembling-hand concepts, and is phrased in terms of a particular definition of trembling-hand perfection, we shall begin by describing eight possible formulations of trembling-hand perfection, and the relations between them.

We shall work in the context of a normal form game $[N, \Sigma, h]$. This is extended to a game $[N, M, H]$ by the use of mixed strategies; $M = \times_{i \in N} M_i$ is the space of mixed strategies; each M_i is the set of Σ_i -valued random variables, or probability distributions over Σ_i . There are situations where one or the other interpretation is to be preferred, but we shall blur the distinction here. H is the expected payoff function. The generic member of $\Sigma = \times_{i \in N} \Sigma_i$ is written $s = [s_1, \dots, s_n]$ while the generic member of $M = \times_{i \in N} M_i$ is written $\sigma = [\sigma_1, \dots, \sigma_n]$, and where we denote by $\sigma_i(s)$ the probability that $s \in \Sigma_i$ will be used if player i employs the mixed strategy σ_i .

4.1 Definition: $\sigma^* \in M$ is Pa iff there exist sequences $\varepsilon_t \subset \mathbb{R}_+$ and $\sigma_t^* \subset M^\circ$ [the interior of $M = \{\sigma \in M: \text{for all } i, s \in \Sigma_i, \sigma_i(s) > 0\}$] such that

- i) $\varepsilon_t > 0$, all t , $\lim_{t \rightarrow \infty} \varepsilon_t = 0$;
- ii) $\lim_{t \rightarrow \infty} \sigma_t^* = \sigma^*$; and
- iii) for all $t, j \in N$, $s_j \in \Sigma_j$, if $[\sigma^* | s_j]$ denotes the strategy n -tuple $[\sigma_1^*, \dots, \sigma_{j-1}^*, s_j, \sigma_{j+1}^*, \dots, \sigma_n^*]$, then $H_j([\sigma^* | s_j]) < \max_{s \in \Sigma_j} \{H_j([\sigma^* | s])\}$ implies $\sigma_t^*(s_j) \in (0, \varepsilon_t)$.

4.2 Definition: $\sigma^* \in M$ is Pb iff for every sequence $\varepsilon_t \subset R_+$ [s.t. $\varepsilon_t \rightarrow 0$] there exists a sequence $\sigma_t^* \subset M^\circ$ s.t. i-iii above obtain.

4.3 Definition: $\sigma^* \in M$ is Pc iff there exists a sequence $\bar{\eta}_t$ of maps $\bar{\eta}_t: \Sigma \rightarrow (0,1)$ and a sequence $\sigma_t^* \subset M^\circ$ s.t.

iv) for all $i \in N$, t and $s \in \Sigma_i$, $\bar{\eta}_t(s) > 0$, and $\lim_{t \rightarrow \infty} \bar{\eta}_t(s) = 0$;

ii) [as above]; and

v) for each $t, j \in N$, $s_j \in \Sigma_j$, $H_j([\sigma^*|s_j]) < \max_{s \in \Sigma_j} \{H_j([\sigma^*|s])\}$
implies $\sigma_t^*(s_j) \in (0, \bar{\eta}_t(s_j))$

4.4 Definition: $\sigma^* \in M$ is Pd iff for all sequences $\bar{\eta}_t$ of maps

$\bar{\eta}_t: \Sigma \rightarrow (0,1)$ s.t. for all j , $s \in \Sigma_j$, $\lim_{t \rightarrow \infty} \bar{\eta}_t(s) = 0$, there exists a sequence $\sigma_t^* \subset M^\circ$ such that ii) and v) are satisfied.

4.5 Definition: $\sigma^* \in M$ is Pe iff there exists a sequence $\varepsilon_t \subset R_+$ and a sequence $\sigma_t^* \subset M^\circ$ s.t. i), ii), and

vi) if $M(\varepsilon_t)$ denotes $\{\sigma \in M: \text{for all } i, s \in \Sigma_i, \sigma(s) \geq \varepsilon_t\}$,

then for each $i \in N$, and each $\tau_i \in M_i(\varepsilon_t)$,

a) $\sigma_t^* \in M(\varepsilon_t)$, and

b) $H_i(\sigma^*) \geq H_i([\sigma^*|\tau_i])$

[in words, for each t , σ_t^* is a Nash Equilibrium of the game $[N, M(\varepsilon_t), H]$]

4.6 Definition: $\sigma^* \in M$ is Pf iff for every sequence $\varepsilon_t \subset R_+$ s.t.

i) is satisfied, there exists a sequence $\sigma_t^* \subset M^\circ$ s.t. ii) and vi) are satisfied.

4.7 Definition: $\sigma^* \in M$ is Pg iff there exists a sequence of maps $\bar{\eta}_t$, where $\bar{\eta}_t: \Sigma \rightarrow (0,1)$ and a sequence $\sigma_t^* \subset M^\circ$ s.t. ii), iv), and

vii) if $M(\bar{\eta}_t)$ denotes $\{\sigma \in M: \text{for all } i, s \in \Sigma_i, \sigma(s) \geq \bar{\eta}_t(s)\}$

then $\sigma_t^* \in M(\bar{\eta}_t)$, and for all $i \in N$, and for each

$\tau_i \in M_i(\bar{\eta}_t)$, $H_i(\sigma^*) \geq H_i([\sigma^*|\tau_i])$.

4.8 Definition: $\sigma^* \in M$ is Ph iff for every sequence $\bar{\eta}_t$ of maps satisfying iv), there exists a sequence $\sigma_t^* \in M^\circ$ s.t. ii) and vii) are satisfied.

Now that we have all these definitions, it is of some interest to know what the relations between them are. Essentially $Pa = Pe$; $Pb = Pf$; $Pc = Pg$; and $Pd = Ph$. Moreover, $Pd \neq Pb \neq Pa \neq Pc$.

When we consider notions of properness, the crucial element is a metric on the space of strategies, so that we can be precise about what constitutes a small mistake. For example, in Myerson's original definition, this metric can be derived from the payoff function. The various pure strategies, σ_i of player i are ranked according to $H_i(\sigma_i, \mu_{(i)})$, where μ denotes the intended n -tuple of mixed strategies. We can define a metric $d: \Sigma_i \times M_{(i)} \rightarrow R_+$, which we write $d(\sigma_i | \mu_{(i)})$ to indicate that it is a metric on Σ_i , and that the dependence on $\mu_{(i)}$ is parametric, as follows:

$$d(\sigma_i | \mu_{(i)}) = [\max_{\tau \in \Sigma_i} H_i(\tau, \mu_{(i)})] - H_i(\sigma_i, \mu_{(i)})$$

The "distance" being measured is the distance of the given pure strategy from the best response: the "size of the mistake". If we are given any other "distance" $d^*: \Sigma_i \rightarrow R_+$ which measures how far a particular pure strategy is from the best response, we can construct a properness concept by requiring that the probability with which a given pure strategy is used be dominated by $f \circ d^*(\sigma)$, where f is a monotonic decreasing function $f: R_+ \rightarrow (0,1]$, with $f(0) = 1$. If we consider a sequence of such functions, f^t , which converge (in some suitable sense) to the function f^* given by:

$$f^*(d) = \begin{cases} 0 & \text{iff } d > 0 \\ 1 & \text{if } d = 0 \end{cases}$$

then we have the makings of a proper properness notion.

4.9 Definition: $d_i^*: \Sigma_i \times M_{(i)} \rightarrow R_+$ is a proper distance iff

$$[\sigma \in \Sigma_i, \sigma \in \operatorname{argmax}_{t \in \Sigma_i} H_i(t, \mu_{(i)})] \Leftrightarrow [d_i^*(\sigma, \mu_{(i)}) = 0], \text{ and}$$

$\underline{d}_i = \sup\{d_i^*(\tau, \mu): \tau \in \Sigma_i, \mu \in M_{(i)}\}$ is attained and is finite, each i .

4.10 Definition: A family of functions $\{f^t: t = 1, \dots, f^t: R_+ \rightarrow (0, 1]\}$

is said to be admissible iff

- i) $f^t(0) = 1$, all t ;
- ii) f^t is monotone decreasing, all t ;
- iii) $t > t' \Leftrightarrow f^t(d) < f^{t'}(d)$ for all $d > 0$;
- iv) $\lim_{t \rightarrow \infty} f^t = f^*$ [see above]

The family is said to be strongly admissible if the f^t are strictly monotone decreasing for all t . Obviously, the admissibility of a family depends on the convergence notion used in 4.10.iv.

4.11 Definition: Let a proper distance, d , and a member f^t of an admissible family of functions be given. An n -tuple, $\mu \in M^\circ$ of mixed strategies is said to comprise an f^t -proper equilibrium w.r.t. d iff for each i , and each $\sigma_i \in \Sigma_i$,

$$f^t(\underline{d}_i) \leq \mu_i(\sigma_i) \leq f^t(d_i(\sigma_i, \mu_{(i)}))$$

An n -tuple $\mu^* \in M$ of mixed strategies is said to be a proper equilibrium w.r.t. a proper distance d iff there exists an admissible family of

functions $\{f^t\}$ and a sequence $\{\mu^t\} \subset M^o$ of n -tuples of mixed strategies with the following properties:

- i) for each t , μ^t is an f^t -proper equilibrium w.r.t. d ; and
- ii) $\lim_{t \rightarrow \infty} \mu^t = \mu^*$.

4.12 Remark: It is clear that a proper equilibrium w.r.t. the distance $d(\sigma_i | \mu_{(-i)})$ defined on the previous page is precisely a proper equilibrium in the sense of Def. 3.11. The following theorem shows that this is not a vacuous concept.

4.13 Theorem: Let $[N, \Sigma, h]$ be a normal form game where the pure strategy spaces are finite. Then for any proper distance d there exists a proper equilibrium w.r.t. d .

Proof: We begin by showing that for each t there exists an f^t -proper equilibrium w.r.t. d . Let $\underline{M}_1^t = \{\mu_1 \in M_1 : \mu_1(\sigma) \geq f^t(\underline{d}_1), \text{ all } \sigma \in \Sigma_1\}$. It is immediate that $\underline{M}^t = \prod_{i \in N} \underline{M}_i^t$ is a compact subset of M^o , and is non-empty. We now define a correspondence $\Psi: \underline{M}^t \rightarrow \underline{M}^t$ by

$$\Psi_i(\mu) = \{\mu_i \in \underline{M}_i^t : \mu_i(\sigma) \in [f^t(\underline{d}_i), f^t(d_i(\sigma, \mu_{(-i)}))], \text{ all } \sigma \in \Sigma_i\}$$

Since the pure strategy spaces are finite, this mapping is upper semicontinuous, and, being the set of solutions to a finite set of linear inequalities, is closed- and convex-valued. Hence the mapping Ψ has a fixed point, which is obviously an f^t -proper equilibrium w.r.t. d . The final step is that, since M is a compact set, the sequence of fixed points has a limit, which is a proper equilibrium w.r.t. d .

4.14 Remark: We can relax the finiteness condition, as long as we restrict attention to distances and admissible families which preserve the uppersemicontinuity and convex-valuedness of the correspondence.

It is also evident that the proper equilibrium w.r.t. any proper distance d form a subset of the set of trembling-hand perfect equilibria. Finally, if we are given two distances d and d' with the property that d is a monotone increasing transformation of d' , then the set of proper equilibria w.r.t. d coincides with the set of proper equilibria w.r.t. d' . To see this, if $g \circ d = d'$ and if f^t is an admissible family that works for d' , then $f^{t \circ g}$ is an admissible family that works for d .

B. Relations between the various notions defined so far

We begin by recapitulating the perfectness concepts used up to now.

P1: subgame perfection for single-vertex subgames (Selten)

P2: subgame perfection for all subgames (Cave)

P3: trembling-hand perfection in the normal form (Selten)

P4: properness (Myerson)

P5: perfection (Selten)

We further remark that P3 perfection has four different definitions, depending on whether we require convergence for one or all sequences of minimum weights, and whether these weights are uniform or can be varied for each player and each strategy. However, for the purposes of the present classification we shall stick to the definition given in Def. 3.9 above.

4.15 Theorem: The following implications characterize the various perfection notions given above.

$$P6 \Rightarrow P5 \Rightarrow P3 \Rightarrow P2 \Rightarrow P1$$

$$P4 \Rightarrow P3 \Rightarrow P2 \Rightarrow P1$$


None of the converse implications holds.

Proof: We begin by remarking that P6 is properness in the agent-normal form of an extensive form game of perfect recall. Thus it is clear that $P6 \Rightarrow P5$: the reasoning is that used in establishing $P4 \Rightarrow P3$. It is also obvious that $P2 \Rightarrow P1$, since the former requires local equilibrium at a strictly larger set of subgames. The fact that P5 is the same as P3 applied to the agent-normal form shows that $P5 \Rightarrow P3$. Finally, the same argument Selten used to establish $P3 \Rightarrow P1$ shows that $P3 \Rightarrow P2$; P2 invokes a richer set of subgames, but the strategic choices are still richer in the normal form, where mixed strategies can be employed.

To show that the converses are false, we present a series of examples.


4.16 Examples:

A. $P1 \not\Rightarrow P2$: consider the following game in extensive form: $(0, 10)$



There are equilibria at which player 1 chooses "top" at I_1^1 and "bottom" at I_1^1 , if player 2 threatens to play "bottom" at I_2^2 . This threat would not be considered credible in the subgame defined by I_2^1 but since this is not considered a subgame in the sense of P1, perfection for these equilibria is possible. According to P2, however, the only perfect equilibrium involves player 2 playing "top" at I_2^1 , while player 1 plays "top" at I_1^2 and "bottom" at I_1^1 , giving each player a payoff of $4-3p$.

B. $P2 \neq P5$: consider the following three-player game



The normal form is given by:

	L_3		R_3	
	L_2	R_2	L_2	R_2
Tt	(1, 3, 0)	(2, 0, 0)	(1, 3, 0)	(2, 0, 0)
Tb	(1, 3, 0)	(0, 0, 5)	(1, 3, 0)	(4, 4, 0)
Bt	(0, 0, 0)	(0, 0, 0)	(3, 0, 3)	(3, 0, 3)
Bb	(0, 0, 0)	(0, 0, 0)	(3, 0, 3)	(3, 0, 3)

The agent normal form is described as follows: the agents of player 1 choose the rows as indicated, while 2 and 3 choose the columns.

	L_3		R_3	
	L_2	R_2	L_2	R_2
T	(1, 1, 3, 0)	(2, 2, 0, 0)	(1, 1, 3, 0)	(2, 2, 0, 0)
t				
B	(0, 0, 0, 0)	(0, 0, 0, 0)	(3, 3, 0, 3)	(3, 3, 0, 3)
T	(1, 1, 3, 0)	(0, 0, 0, 5)	(1, 1, 3, 0)	(4, 4, 4, 0)
b				
B	(0, 0, 0, 0)	(0, 0, 0, 0)	(3, 3, 0, 3)	(3, 3, 0, 3)

Since this game has no subgames, either in the P1 or P2 sense, it follows that the sets of equilibria, P1 perfect equilibria, and P2

perfect equilibria are all identical. Since this is a fairly complicated game, we shall focus on particular pure-strategy equilibria. In the normal form, the strategies Bt and Bb are identical, so when we analyze this game we can specify normal-form strategies as $[(p_1, p_2), q, r]$, where p_1 is the probability of Tt, p_2 is the probability of Tb, q is the probability of L_2 , and r is the probability of L_3 . Thus the payoff functions can be written:

$$H_1() = p_1[3r-q-1] + p_2[4qr-3q-r+1] + 3[1-r]$$

$$H_2() = q[3p_1-p_2+4p_2r] + 4p_2[1-r]$$

$$H_3() = r[8p_2-5p_2q+3p_1-3] + 3[1-p_1-p_2]$$

The best response functions are given by:

$$\beta_1(q, r) = \begin{cases} (1, 0) & \text{if } x > \max[0, y] \\ (0, 1) & \text{if } y > \max[0, x] \\ (0, 0) & \text{if } 0 > \max[x, y] \\ (p, 1-p) & \text{if } x = y > 0 \\ (p_1, p_2) & \text{if } x = y = 0 \end{cases}$$

where $x = 3r-q-1$, and $y = 4qr-3q-r+1$

$$\beta_2(p, r) = \begin{cases} 1 & \text{if } 3p_1 > p_2(1-4r) \\ q & \text{if } 3p_1 = p_2(1-4r) \\ 0 & \text{if } 3p_1 < p_2(1-4r) \end{cases}$$

$$\beta_3(p, q) = \begin{cases} 1 & \text{if } 3(1-p_1) > p_2(8-5q) \\ r & \text{if } 3(1-p_1) = p_2(8-5q) \\ 0 & \text{if } 3(1-p_1) < p_2(8-5q) \end{cases}$$

We observe that there are four pure-strategy equilibria: $[(1,0),1,1]$; $[(0,1),1,1]$; $[(0,0),0,0]$; and $[(0,0),1,0]$. All of these are P1 and P2 perfect. To verify that $[(0,0),1,0]$ is P3 perfect, we make use of the following:

4.17 Lemma: If μ is a P3 perfect equilibrium and if $\{\mu^t\}$ is a sequence of n-tuples of perturbed-game equilibria that establishes the perfection of μ , then for each i there exists a $T(i)$ with the property that, for all $t \geq T(i)$, $\mu_i \in \Delta_i(M_i, \mu_{(-i)}^t)$. In other words, if player i is allowed to choose any mixed best response to the perturbed equilibrium of games sufficiently far out in the sequence, his choice will be the same as at the limit equilibrium; however, it is certainly not true that this choice will necessarily be in equilibrium with the choices of the other players in the perturbed game.

Proof: Linearity of payoff function and convergence of $\{\mu^t\}$.

Returning to our example, we shall construct a sequence of completely mixed minimum probability assignments by stipulating that v^t is a sequence of small positive constants depending only on t . That is, every pure strategy not used in $[(0,0),1,0]$ is used with probability v^t . We assume the v^t are sufficiently small and converge to 0. Thus we have a sequence of n-tuples of the form $[(v^t, v^t), 1-v^t, v^t]$, to which it is obvious that the pure strategies $(0,0) = (p_1, p_2)$, $1 = q$, and $0 = r$ are best replies. Therefore, $[(0,0),1,0]$ is P3 perfect.

However, it is not P5 perfect. This can be seen from the agent-normal form of the game: consider any sequence of completely mixed strategies converging to $[0,0,1,0]$, which is the agent-normal version of $[(0,0),1,0]$. When the agent of player 1 at I_1^2 comes to move he

knows that both the other agent of player 1 and player 2 made mistakes. Thus, he faces the following game with player 3: (we show only their payoffs)

	L_3	R_3
t	(2,0)	(2,0)
b	(0,5)	(4,4)

It is now immediate that if the probability that player 3 will play R_3 (which is what the equilibrium calls for) is sufficiently high, the best response for the agent of player 1 at I_1^2 is "b", instead of the "t" called for at the equilibrium in question. We have therefore shown that $P2 \neq P5$, and also that $P3 \neq P5$.

Examples showing that the other reverse implications fail can be found in referenced works by Myerson and Selten.

4.18 Remarks: In games of perfect information, all the above concepts coincide. In games of effectively perfect information and perfect recall, $P5$ is equivalent to $P2$, where effectively perfect information means: if $i, j \neq 0$ and $i \neq j$, for any $I_i \in \Xi_i$, $I_j \in \Xi_j$,

$$[I_i \rightarrow I_j] \Rightarrow [I_j \subset D(I_i, k), \text{ some } k].$$

In words, this says that each player can remember all previous moves of other personal players, although he need not be able to recall his own past moves (perfect recall) or those of the chance player. Additionally, if a game has perfect recall it follows that choices can be made independently at each information set. One example of such a game is the so-called standard supergame, which consists of a denumerably infinite repetition of a normal-form game. Each player is allowed, at

each play of the stage game, to base his move on the entire previous history of the game, thus making the game one of perfect recall. However, since the stage game is a normal-form game, the players all move simultaneously at each stage, so the supergame fails to have perfect information. However, the property of perfect recall here is sufficient to guarantee that each agent of a given player has the same information in the game as a whole as he does in the subgame that starts with his actions. Thus, for standard supergames all the perfectness notions coincide.

C. Extension of P5 perfection to games without perfect recall

It will be recalled that none of the perfectness concepts short of P5 required perfect recall; the role it plays in P5 is to guarantee that choices can be made independently at each information set. In games which fail to have perfect recall, we can still retain the intuition behind P5 perfection as long as we identify a canonical set of independent decision problems for each agent. In this section we provide two ways of arriving at this canonical form.

I. The Partition Method

Recall that if $[N, K, \Xi, h, p]$ is a game in extensive form, Ξ is a partition of K into $n+2$ collections of subsets, denoted $\Xi_0, \Xi_1, \dots, \Xi_n, W$. W is the set of endpoints, and $M(W)$ denotes the set of all probability distributions on W that can be induced by the use of mixed strategies. We now limit attention to player i 's strategies. We shall denote by Π a partition of Ξ_i .

4.19 Definition: Let Π be a partition of Ξ_i , and $v \in M_i$ a mixed strategy of player i . We say that v is a Π -strategy iff it can be represented

as $v(p_1), \dots$, where the p_t are elements of Π , and where each $v(p_t)$ is a random variable taking values in $C(p_t)$, the space of $|p_t|$ -tuples of pure moves, one for each information set in p_t . In other words, a Π -strategy is a mixed strategy which may be mixed within elements of Π , but which is constrained to be independent between such elements.

4.20 Examples: Π is the coarse (one-element) partition; a Π -strategy is a mixed strategy; $\bar{\Pi}$ is the fine partition (each information set is an element; a $\bar{\Pi}$ -strategy is a behavioral strategy.

4.21 Definition: If $w(\mu)$ denotes that member of $M(W)$ which results when the strategy n -tuple $\mu \in M$ is employed, and if $S_i(\Pi)$ denotes the space of player i 's Π -strategies, we define

$$\Phi(\Pi) = \{w(v_i, \mu_{(i)}): v_i \in S_i(\Pi), \mu_{(i)} \in M_{(i)} = S_{(i)}(\underline{\Pi})\}$$

i.e., that subset of $M(W)$ which represents the probability distributions attainable when player i is restricted to the use of Π -strategies.

4.22 Proposition: $\Pi \subset \Pi' \Rightarrow \Phi(\Pi) \supset \Phi(\Pi')$

Proof: Since Π is coarser than Π' , the independence requirements on Π' -strategies are weaker than those on Π -strategies: every Π -strategy is already a Π' -strategy.

4.23 Proposition: $\Phi(\Pi) = \Phi(\Pi') = \Phi(\underline{\Pi}) \Rightarrow \Phi(\Pi) = \Phi(\Pi \wedge \Pi')^{1/}$

Proof: Recall that a necessary and sufficient condition for the hypothesis of the theorem is that moves in any element of Π or Π' have perfect recall w.r.t. moves in other elements. In precise terms, if $u \in p \in \Pi$, and $v \in p^* \neq p$, $p^* \in \Pi$, then either

- i) there exists no $w \in W$ s.t. $\bar{w} \cap u \neq \emptyset \neq \bar{w} \cap v$; or
- ii) if, w.l.o.g., $D(u) \cap v \neq \emptyset$, then $v \subset D(u, k)$, some k .

Thus, the proposition will be proven if we can show that this property is inherited by the meet. However, for all $p \neq p'$, where p and p' are elements of $\Pi \wedge \Pi'$, either

iii) there exist $p_1 \neq p_2$, elements of Π , s.t. $u \in p_1, v \in p_2$; or

iv) there exist $p'_1 \neq p'_2$, elements of Π' , s.t. $u \in p'_1, v \in p'_2$.

Q.E.D.

4.24 Definition: $\Pi^* = \wedge \{\Pi: \Phi(\Pi) = \Phi(\underline{\Pi})\}$

By the previous proposition, Π^* is well-defined; it is the finest partition whose admissible strategies are equivalent to the use of mixed strategies, and is equal to Π for the case where the player in question has perfect recall. The elements of this partition are precisely the independent decision situations faced by the player, and will serve as the basis for our extension of P5 perfection.

II. The Signalling Method

In an early paper, Thompson (1953) defined the notion of a signalling strategy for games without perfect recall. The basic idea is that correlation is required, but the extent of this correlation can be limited. In other words, a player's knowledge of the strategy that he employed at an earlier information set can be used to reduce his ignorance at a later date.

4.25 Definition: Let $u \in \Xi_i$ be an information set. u is a signalling information set iff there exists $v \in \Xi_i$ s.t.

i) $v \cap D(u, k) \neq \emptyset$ for some alternative k , and

ii) $v \cap D(u) \not\subset D(u, k)$.

The set of signalling information sets is denoted S_i . For each $u \in S_i$, let $E(u)$ denote the set of v satisfying i) and ii) above.

4.26 Definition: A pure signalling strategy for player i is a choice of a move for each $u \in S_i$. A mixed signalling strategy is a random variable taking values in the space of pure signalling strategies.

4.27 Definition: Let s_i be a pure signalling strategy. $K(s_i)$ is the game tree obtained from K when the choices given by s_i are assigned to each $u \in S_i$, and these information sets are removed from the tree.

$\Xi_i(s_i)$ denotes the restriction of Ξ_i to $K(s_i)$. The information sets in $\Xi_i(s_i)$ will be generically denoted u' . A random variable which takes values in the space of choices at $u' \in \Xi_i(s_i)$ is called a local behavioral strategy at u' : if β_i consists of a local behavioral strategy, $\beta_i(u')$, for each $u' \in \Xi_i(s_i)$, then β_i is called a behavioral strategy associated with s_i . The pair (σ_i, β_i) , where σ_i is referred to as a composite strategy.

4.28 Theorem (Thompson, 1953, Prop. 3): To every mixed strategy there corresponds an equivalent mixed strategy.

4.29 Definition: $u R v$ if either $u \in E(v)$ or $v \in E(u)$; $u \sim v$ iff there exists a chain $u = u_1, \dots, u_m = v$ s.t. $u_k R u_{k+1}$ for all $k = 1, \dots, m-1$. Clearly, $\sim$ is an equivalence relation.

4.30 Definition: $\Pi_* = \Xi_i / \sim$

4.31 Remarks: Since choices at u are correlated with choices at members of $E(u)$, the equivalence relation defined above associates information sets which are necessarily correlated. The elements of the partition Π_* correspond to "agents" defined as follows:

4.32 Definition: Let Π be a partition of Ξ_i into sets $p_1, \dots$. If we denote by H_t the set of moves at p_t ($|p_t|$ -tuples of moves; one for each

$u \in p_t$), we can ascertain whether or not moves at H_t satisfy the condition of perfect recall with respect to moves at $H_{t'}$, $t \neq t'$. If they do, for all H_t , we call Π an agent partition of Ξ_1 . This is not quite enough for our purposes, however, since we need to obtain the finest such partition. However, by examining the proof of 4.23, it is clear that such a finest partition exists, and is equal to Π^* . Moreover, it is clear that

4.33 Proposition: $\Pi^* = \Pi_*$.

The remaining step, now that we have constructed the canonical form of the game, is to extend Selten's notion of (P5) perfection to this canonical form. However, since the canonical form describes a game with perfect recall, this can be done either by requiring that each agent play a local best response at H_t , or by looking at trembling-hand perfection in the "agent normal form" obtained using the agents defined in 4.32.

We remark that the same existence proofs used by Selten to show that P5 equilibria exist work here, so that this is not a vacuous concept. In our next note, we shall apply this concept to some examples of games without perfect recall, in particular, to supergames played under memory limitations.

Footnote:

¹That this result fails in case $\phi(\Pi) \neq \phi(\underline{\Pi})$ is shown in the following example, suggested by Charles Blair.

There are three 2-element partitions:

$$\Pi^1 = (u_1 u_2)(u_3); \Pi^2 = (u_1 u_3)(u_2); \Pi^3 = (u_2 u_3)(u_1)$$

We shall show that $\phi(\Pi^2) = \phi(\Pi^3) \neq \phi(\Pi^2 \wedge \Pi^3) = \phi(\underline{\Pi})$

a) $\phi(\Pi^2) = \phi(\Pi^3)$ a Π^2 -strategy is

$$\begin{aligned} (q_1, q_2, q_3, r) \text{ where } & q_1 = \text{pr}(T_1 T_3) \\ & q_2 = \text{pr}(T_1 B_3) \\ & q_3 = \text{pr}(B_1 T_3) \\ & r = \text{pr}(T_2) \end{aligned}$$

analogously, a Π^3 -strategy is $(q_1^1, q_2^1, q_3^1, r^1)$

$$\begin{aligned} \text{where } & q_1^1 = \text{pr}(T_2 T_3) \\ & q_2^1 = \text{pr}(T_2 B_3) \\ & q_3^1 = \text{pr}(B_2 T_3) \\ & r^1 = \text{pr}(T_1) \end{aligned}$$

The realization probabilities for the various endpoints are:

	Π^2	Π^3
pr(a)	$(q_1 + q_2)r$	$(q_1^1 + q_2^1)r^1$
pr(b)	$q_1(1-r)$	$q_3^1 r^1$
pr(c)	$q_2(1-r)$	$(1 - q_1^1 - q_2^1 - q_3^1)r^1$
pr(d)	$q_3 r$	$q_1^1(1-r^1)$
pr(e)	$(1 - q_1 - q_2 - q_3)r$	$q_2^1(1-r^1)$
pr(f)	$(1 - q_1 - q_2 - q_3)(1-r)$	$(1 - q_1^1 - q_2^1)(1-r^1)$

equating the realization probabilities for each endpoint, we see that

there is a unique solution:

$$\begin{aligned} r &= \frac{q_1^1 + q_2^1}{q_1 + q_2} & q_2 &= r^1 \left(\frac{1 - q_1^1 - q_2^1 - q_3^1}{1 - q_1^1 - q_2^1} \right) \\ q_1 &= \frac{r^1 q_3^1}{1 - q_1^1 - q_2^1} & q_3 &= (1 - r^1) \left(\frac{q_1^1}{q_1^1 + q_2^1} \right) \end{aligned}$$

which shows equivalence of Π^2 and Π^3

$$b) \quad \phi(\Pi^2) \neq \phi(\Pi^2 \wedge \Pi^3): \quad \Pi^2 \wedge \Pi^3 = (u_1)(u_2)(u_3)$$

so if $\phi(\Pi^2) = \phi(\Pi^2 \wedge \Pi^3)$ we would be able to find (q_1, q_2, q_3, r) s.t.

$$\begin{aligned} (q_1 + q_2)r &= \alpha_1 & q_3r &= \alpha_4 \\ q_1(1-r) &= \alpha_2 & (1-q_1-q_2-q_3)r &= \alpha_5 \\ q_2(1-r) &= \alpha_3 & (1-q_1-q_2-q_3)(1-r) &= 1-\alpha_1-\alpha_2-\alpha_3-\alpha_4-\alpha_5 \end{aligned}$$

$$\text{for any } \alpha_i \text{ s.t. } \sum_{i=1}^5 \alpha_i \leq 1$$

However, it is easy to see that this can only work when, for example,

$$\frac{\alpha_1}{\alpha_2 + \alpha_3} = \frac{\alpha_5}{1 - \alpha_1 - \alpha_2 - \alpha_3 - \alpha_4 - \alpha_5} : \text{ an additional constraint on the } \alpha_i \text{'s.}$$

References:

- Cave, J.A.K., Equilibrium Behavior in Infinitely-repeated Games, Ph.D. dissertation, Stanford, 1980, esp. Ch. V.
- Kuhn, H., "Extensive Games and the Problem of Information" in Contributions to the Theory of Games, vol. II, Princeton University Press, 1953.
- Luce, R. and H. Raiffa, Games and Decisions, John Wiley, Inc., 1957.
- Myerson, R., "Refinements of the Nash Equilibrium Concept," Northwestern University Discussion Paper 295, August, 1977.
- Selten, R., "A Model of Oligopoly where Four are Few and Six are Many," Int. J. Game Theory 2, 1973, pp. 141-201.
- Selten, R., "A Reexamination of the Perfectness Concept for Games in Extensive Form," Int. J. Game Theory 4, 1975, pp. 25-55.
- Thompson, G., "Signalling Strategies in N-person Games," in Contributions to the Theory of Games, vol. II, Princeton University Press, 1953.

HECKMAN
BINDERY INC.



JUN 95

Sound-To-Please® N. MANCHESTER,
INDIANA 46962

UNIVERSITY OF ILLINOIS-URBANA



3 0112 060296214